

# Laser: Efficient Language-Guided Segmentation in Neural Radiance Fields

Xingyu Miao, Haoran Duan\*<sup>✉</sup> *Member, IEEE*, Yang Bai, Tejal Shah, Jun Song, Yang Long<sup>✉</sup>, *Senior Member, IEEE*, Rajiv Ranjan, *Fellow, IEEE*, and Ling Shao, *Fellow, IEEE*

**Abstract**—In this work, we propose a method that leverages CLIP feature distillation, achieving efficient 3D segmentation through language guidance. Unlike previous methods that rely on multi-scale CLIP features and are limited by processing speed and storage requirements, our approach aims to streamline the workflow by directly and effectively distilling dense CLIP features, thereby achieving precise segmentation of 3D scenes using text. To achieve this, we introduce an adapter module and mitigate the noise issue in the dense CLIP feature distillation process through a self-cross-training strategy. Moreover, to enhance the accuracy of segmentation edges, this work presents a low-rank transient query attention mechanism. To ensure the consistency of segmentation for similar colors under different viewpoints, we convert the segmentation task into a classification task through label volume, which significantly improves the consistency of segmentation in color-similar areas. We also propose a simplified text augmentation strategy to alleviate the issue of ambiguity in the correspondence between CLIP features and text. Extensive experimental results show that our method surpasses current state-of-the-art technologies in both training speed and performance. Our code is available on: <https://github.com/xingy038/Laser.git>.

**Index Terms**—3D segmentation, CLIP, NeRF.

## 1 INTRODUCTION

ADVANCEMENTS in Neural Radiance Fields (NeRFs) [1] and 3D Gaussian splatting (3D GS) [2] have driven progress in novel view synthesis of 3D scenes in recent years. However, these models often focus primarily on geometry and appearance modeling with limited consideration of semantic information, thereby affecting the overall understanding of 3D scenes. Comprehensive understanding of 3D scenes is crucial for applications such as robot navigation, or autonomous driving. Several endeavors [3], [4] have aimed to enrich 3D scenes with multi-view semantic information. Despite these advancements, the scarcity of adequately annotated 3D scene datasets continues to pose a significant hurdle. An effective approach to address this issue involves adopting open vocabulary segmentation, which facilitates the identification and classification of diverse semantic regions and objects within these scenes. However, this strategy introduces its own set of substantial challenges, chiefly the requirement for an in-depth un-

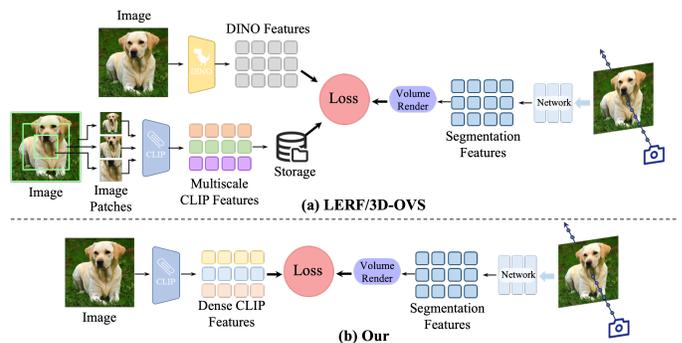


Fig. 1. **Workflows** of existing methods and Ours: (a) The core of the LERF/3D-OVS process initially adopts a cutting strategy that subdivides the training images into patches of different sizes. These patches are then fed into the CLIP encoder to extract multi-scale CLIP features, which are subsequently saved. At the same time, the original image is also input into DINO to extract DINO features. Afterwards, these multi-scale CLIP features and DINO features participate together in the segmentation branch optimization process of NeRF. Such a process demands significantly high computational and storage resources. (b) Compared to previous methods, our process only requires the input of training images into a modified CLIP encoder, after which it can predict dense CLIP features. Utilizing these features, the segmentation branch of NeRF can be optimized efficiently, while significantly reducing the consumption of computational and storage resources.

- Xingyu Miao and Yang Long are with the Department of Computer Science, Durham University, UK. E-mail: xingyu.miao@durham.ac.uk; yang.long@ieee.org.
- Yang Bai is with the Institute of High Performance Computing (IHPC), ASTAR, Singapore. E-mail: bai\_yang@ihpc.a-star.edu.sg
- Jun Song is with School of Computer Science, China University of Geosciences, Wuhan, 430074, P. R. China. E-mail: songjun@cug.edu.cn
- Rajiv Ranjan, Haoran Duan and Tejal Shah are with the School of Computing, Newcastle University, UK. E-mail: rranjans@gmail.com, haoran.duan@ieee.org, tejal.shah@newcastle.ac.uk
- Ling Shao is with the UCAS-Terminus AI Lab, University of Chinese Academy of Sciences, Beijing 100049, China. E-mail: ling.shao@ieee.org.
- Yang Long and Haoran Duan are the corresponding author.
- Xingyu Miao\* and Haoran Duan\* have equal contribution.

Manuscript submitted March 15, 2024;

derstanding of both natural language processing and the accurate representation of objects in 3D environments.

Recent work [5]–[7] distills features from existing visual language models, such as CLIP [8], into 3D scenes for performing open-vocabulary query tasks. This task aims to enable human users to interact with large language models (LLM) [8]–[11] for scene editing intuitively. In comparison to traditional semantic labeling methods, the language features

of visual language models provide a more comprehensive semantic understanding, especially encompassing objects in the long-tail distribution, making them more applicable in practical scenarios. However, accurately integrating language embeddings into the current 3D scene representation while maintaining efficiency and visual quality poses a few significant challenges. One challenge is aligning the learned 3D points and CLIP embeddings during rendering since CLIP embeddings are represented at the image level rather than the pixel level. To address this issue, LERF [5] and 3D-OVS [6] employ a patch cropping approach to obtain CLIP embeddings at different scales. However, this approach introduces a challenge related to clarity or precision in identifying semantic scales associated with the same 3D position. For instance, if we consider a point on the left fan of a graphics card, different levels of specificity in textual queries can lead to the identification of regions that correspond to varying scales of detail. Specifically, text queries such as "graphics card," "left side of the graphics card," and "fan on the left side of the graphics card" may each highlight areas that are relevant at three distinct levels of granularity. To address this problem, the LERF simply averages the multi-scale features, while 3D-OVS uses a 3D selection volume to choose the most appropriate scale. Despite their innovative approaches, both methods significantly hamper processing efficiency. This inefficiency is evident in the pre-processing phase, where prioritizing the extraction and storage of multi-scale features necessitates a substantial storage capacity—approximately 35GB for the features of each scene. Furthermore, multi-scale patch features bring another issue where they cannot fully include a segmented object, either including other objects in one patch or the objects in the patch are incomplete. Due to these inaccurate CLIP features, the 3D language embedding fields lack clear boundaries and contain a significant amount of noise. To alleviate this issue, these methods extract pre-trained DINO features [12] concurrently during training and align them with 3D segmentation features. However, this approach leads to a further increase in training time and extra information requirements (The workflow as shown in Figure 1 (a)).

Driven by the goal of achieving rapid training speeds without compromising on accuracy in 3D language embedding fields, this paper presents Laser - an efficient end-to-end framework for segmenting 3D neural radiance fields, designed to balance high performance with computational efficiency (The overall workflows of our work as shown in Figure 1 (b)). In order to obtain pixel-level features from CLIP, we draw inspiration from recent advancements in CLIP dense prediction research, such as MaskCLIP [13], CAT-Seg [14], and SCLIP [15]. These methods modify the final pooling layer within the CLIP encoder to derive dense image embeddings. However, these methods introduce significant noise without a mechanism for refinement, potentially resulting in inaccurate predictions. Although aligning these dense features with 3D segmentation features directly can yield a roughly segmented image, there are still significant amounts of noise present. To address this issue, we propose a novel adapter that effectively reduces the noise present in the segmented images through the incorporation of a loss term and a self-cross-training strategy during

training. Although this can alleviate the noise problem and obtain relatively accurate segmentation maps, it cannot guarantee the consistency of different views in the same scene. To maintain segmentation consistency across different views, we introduce label volumes to transform the segmentation problem into a classification task. In this way, we can classify rays within the segmentation volume, thus ensuring consistent segmentation results under different views. However, we observe that two query texts in the same category area may become ambiguous due to variations in lighting and image color. This ambiguity may lead to noise in the same segmentation region (i.e. regions segmented as the same category may include other categories). To alleviate this issue, we propose a novel low-rank transient query attention and a simplified text augmentation strategy to enhance features. In summary, the contributions of our work are as follows:

- We review prior 3D language-guided segmentation methods that primarily target broad modality alignment. To attain finer alignment, we introduce a streamlined framework designed for efficient segmentation through 3D language guidance.
- We introduce a novel adapter module and self-crossing method to reduce the noise impact caused by dense CLIP features.
- We introduce a novel low-rank transient query attention model that significantly reduces the computational complexity required by vanilla attention mechanisms when processing 3D features.
- We carefully propose a simplified text enhancement strategy that significantly reduces the ambiguity of similarity calculation between CLIP features and text features, thereby effectively improving the accuracy of feature matching.

## 2 RELATED WORK

### 2.1 Neural Rendering

Image-based novel view synthesis methods typically employ input images to integrate pixel information, and their performance can be enhanced through multi-view approaches. Light field and Lomography rendering utilize filtering and interpolating ray samples to generate novel views without the need for an explicit geometric model. However, they necessitate capturing and storing numerous views for high-quality rendering [16]–[20]. Many approaches leverage explicit proxy geometry to achieve high-quality rendering with only a few input images [17], [21]–[29]. However, these methods may suffer from untextured regions, highlights, reflections, and repetitive patterns, posing challenges for accurately estimating scene geometry. To address these issues, previous works have explored techniques such as local warps, soft 3D reconstruction, and learning-based approaches.

Recently, neural implicit representation methods like NeRF [1], [30]–[35] have demonstrated significant potential for high-quality rendering. NeRF employs multi-layer perceptrons (MLPs) to implicitly represent continuous scenes, yielding impressive view synthesis results. The advent of NeRF marks substantial progress in the neural network representation of 3D scenes. However, these methods require

training a separate model for each scene, and optimizing these models entails varying training times. Several recent methods have effectively improved rendering and training speed through explicit and hybrid scene representation [36]–[39]. These methods utilize explicit spatial data structures to store explicit scene data [39], [40] or features decoded by micro MLP [36]–[38]. This separates the model’s capacity from its speed, allowing real-time rendering of high-quality images [37]. However, embedding language representations directly onto the neural field of implicit expressions is not an easy task, as a large number of sampling points require extensive memory usage, resulting in a sharp decline in training and rendering performance.

## 2.2 2D Language Segmentation

The goal of 2D language segmentation is to understand arbitrary categories in images with textual descriptions. In the previous works, ZS3Net [41] is a pioneering work that utilizes word embeddings of unseen classes through a generative model to generate pixel-level features. SPNet [42] uses word embeddings (e.g., word2vec [43]) to achieve alignment of semantic and visual features. GroupViT [44] proceeds by grouping segmentation masks directly from text supervision. In the recent methods, they utilize pre-trained CLIP [8] for open-lexical semantic segmentation. LSeg [45] adopts a strategy of aligning pixel embeddings with text embeddings of corresponding semantic classes generated by the CLIP text encoder. Different from pixel-level LSeg, OpenSeg [46] proposes a method to achieve segment-level visual feature alignment with text embeddings through a regional word basis. 3D-OVS [6] is the first work on implicit neural field open vocabulary segmentation. However, this method is essentially trained and tested on a closed vocabulary set. It achieves effective segmentation of 3D scenes by aligning multi-scale CLIP features with 3D segmentation features. On the contrary, to improve efficiency, we adopt the latest CLIP dense prediction method to obtain dense CLIP features [13], avoiding using memory-intensive and computationally slow multi-scale CLIP features.

## 2.3 3D Language Fields

Due to the challenges in accurately identifying 3D regions, incorporating specific semantics into the implicit representation of NeRF based on MLP is extremely difficult. Some methods [4], [47]–[49] add new network branches and attempt to integrate semantic information into NeRF. For instance, Semantic NeRF [50] jointly encodes semantics with appearance and geometric shapes in NeRF to achieve novel semantic view synthesis. Neural Feature Fusion Fields [51] use self-supervised 2D semantic features to distill 3D features. Recent work explores integrating language into 3D scenes for segmentation. For example, DFF [7] adds a branch for language prediction, attempting to segment unseen text labels during training by fitting the feature map of LSeg [45] into NeRF. LERF [5] is the first method to embed CLIP features [8] into NeRF, utilizing powerful CLIP representations for text-based 3D queries. Additionally, they use DINO features [12] to supervise LERF and enhance its performance. 3D-OVS [6] optimizes LERF and achieves open-vocabulary segmentation functionality. However, the

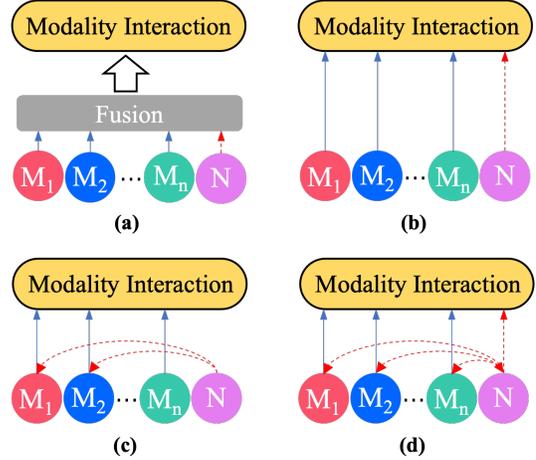


Fig. 2. **Interaction comparison of different modality.** M=modality, N=new modality. (a) Fusion of features from different modalities and then interactive processing. When new modalities are included, retraining is required. (b) Directly interact with features of different modalities. When new modalities are added, they also need to be retrained. (c) Directly interact with the features of different modalities. After introducing a new modality, similar modal features can be used for feature distillation. (d) After introducing new modalities, we not only adopted the feature distillation strategy of similar modalities, but also directly processed the interactive features between different modalities.

low semantic embedding quality of these methods requires the introduction of additional semantic information such as DINO features [12] or SAM features [52], and the time cost during training is high. In addition, with the rise of 3D GS explicit representation, some recent segmentation methods based on 3D GS have also shown excellent results, such as Langsplat [53], LEGaussians [54], and Feature 3dgs [55]. However, these methods still use the data processing method of 3D-OVS. Although the real-time rendering advantage of 3D GS improves the training speed, there is still a lot of redundancy in the preprocessing stage. In contrast, our proposed Laser can provide an efficient 3D language scene representation.

## 3 PROPOSED METHOD

Previous work essentially treated NeRF features as new modalities to be fused with existing multi-modal models, aiming to handle downstream tasks such as segmentation and detection. The integration of this new modality with the multimodal model can be achieved in several different ways as shown in Figure 2. 3D language-guided segmentation methods achieved cross-modal alignment by refining CLIP image features into the segmentation branch of NeRF, i.e. Figure 2 (c). However, these methods adopt a basic and rough alignment strategy and fail to deeply explore and utilize the potential of this alignment. Therefore, they will use post-processing methods for optimization, including pixel-level segmentation strategies and DINO. Our goal is to simplify these post-processing methods and utilize more refined alignment strategies to achieve efficient language-guided segmentation (Figure 2 (d)). This work aims to embed different modalities into the same vector space, so that in this space, semantic level comparison can be performed through a simple dot product operation. Given

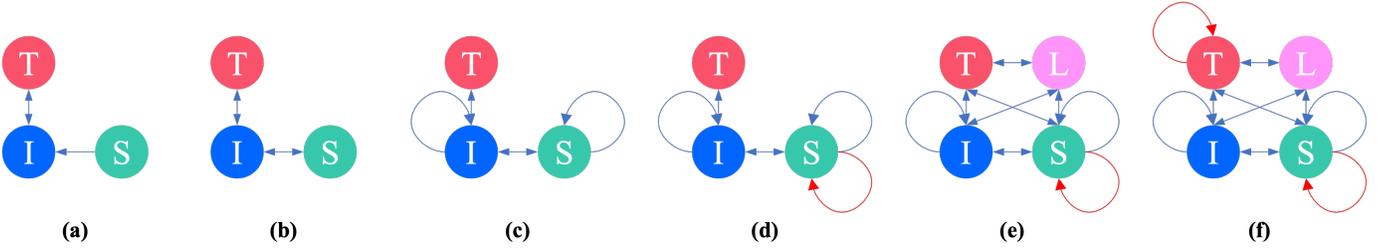


Fig. 3. **Modality Graphs.** T=text feature, I=image feature, S=segmentation feature of NeRF, L=label volume feature of NeRF. (a) Previous methods only distilled the image modality capabilities of CLIP into the segmentation branch of NeRF. (b) and (c) demonstrate our attempt to align the segmentation modality of NeRF with the image modality of CLIP, as discussed in Section 3.1, where we introduced an adapter and a self-cross training strategy. (d) describes the self-enhancement of the NeRF segmentation modality, where in Section 3.2, we proposed a low-rank transient query attention. (e) By incorporating the label modality of NeRF, we achieved bilateral alignment among four modalities, as shown in Section 3.3, introducing label volume and  $\mathcal{L}_{CE}$  loss. (f) pertains to the self-augmentation of the textual modality, where in Section 3.4, we proposed a simplified text augmentation strategy.

that three different modalities are involved: image, text, and NeRF segmentation features, these modal spaces can exist in multiple configurations of interrelationships. Alignment is achieved between text and image modalities, so we can exploit this relationship to implement language-based 3D segmentation guidance, as shown in Figure 3.

Specifically, we propose Laser, a novel language-guided segmentation method reconstructing NeRF. By leveraging multi-view images of a given scene and open vocabulary descriptions for each category to segment the reconstructed NeRFs such that each 3D point is assigned a corresponding category label. To achieve this, we map each 3D point in the scene to CLIP features that represent its semantics. As CLIP only can obtain the image-level features, we use a modified CLIP encoder to derive dense image embeddings and learn an adapter for reducing the noise, as described in Section 3.1. However, variations in lighting and color in images can lead to segmentation ambiguity, resulting in inaccurate correlation values and misclassification. Taking a pair of shoes captured in a photo as an example, the corresponding CLIP features encompass the characteristics of the shoes. If the shoes have an all-black appearance, the shoe label is in a light black shade with a white logo, and the shoes are placed on a table with an all-white color and black stripes, CLIP features may mistakenly classify the shoe label as part of the table. We introduce a novel low-rank transient query attention and a simplified text augmentation strategy to enhance semantic features, as described in Section 3.2 and Section 3.4. In addition, to ensure the consistency of segmentation maps among different views in the same scene, we introduced label volume. By transforming the segmentation problem into a classification task, we assign a corresponding category label to each 3D point, thereby enhancing the consistency of segmentation maps across different views, as described in Section 3.3.

### 3.1 Mapping 3D to Dense Pixel-Level CLIP Features

To improve efficiency, we use the feature maps obtained by the recent CLIP dense prediction method [13], [15] instead of the multi-scale CLIP feature maps used by previous methods [5], [6]. After obtaining the dense pixel-level CLIP features, we follow [5]–[7], introducing an additional branch

to render the CLIP feature. We can render the RGB value and the CLIP feature of each ray  $\mathbf{r}$  using volume rendering [1]:

$$\hat{C}(\mathbf{r}) = \sum_{i=1}^N w_i C_i \in \mathbb{R}^3, \hat{F}(\mathbf{r}) = \sum_{i=1}^N w_i F_i \in \mathbb{R}^D, \quad (1)$$

$$\text{where } w_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right) (1 - \exp(-\sigma_i \delta_i)), \quad (2)$$

where  $\sigma_i$  and  $\delta_i$  refer to the volume density and feature attributes of the sampled point  $i$ , with  $w_i$  indicating the weighting of  $F_i$  within the ray  $\mathbf{r}$ , and  $\delta_i$  denoting the distance between adjacent samples. In every training batch, the supervision loss for a set of rays (denoted as  $\mathcal{R}$ ) is defined as the sum of the L2 distance between the rendered and ground truth RGB values, and the cosine similarities  $\cos(\cdot, \cdot)$  between the rendered features and the dense pixel-level CLIP features  $F(\mathbf{r})$ :

$$\mathcal{L}_s = \sum_{\mathbf{r} \in \mathcal{R}} \left( \left\| \hat{C}(\mathbf{r}) - C(\mathbf{r}) \right\|_2 - \cos \langle \hat{F}(\mathbf{r}), F(\mathbf{r}) \rangle \right). \quad (3)$$

After training, given a set of text descriptions  $\mathcal{T} = \{t_i\}_{i=1}^N$  and a CLIP text encoder  $E_t$ , we can get the classes' text features  $F_t = E_t(\mathcal{T}) \in \mathbb{R}^{N \times D}$ . Then we can calculate the cosine similarity between the rendered CLIP features  $\hat{F}(\mathbf{r})$  and the class' text features  $F_t$  to obtain the segmentation logits  $z(\mathbf{r})$  of the ray  $\mathbf{r}$ :

$$z(\mathbf{r}) = \cos \langle \hat{F}(\mathbf{r}), F_t \rangle \in \mathbb{R}^N. \quad (4)$$

We can then get the segmentation maps of the ray  $s(\mathbf{r}) = \text{argmax}(z(\mathbf{r}))$ . Similarly, we can get the segmentation logits  $z_I(\mathbf{r})$  of dense pixel-level CLIP features:

$$z_I(\mathbf{r}) = \cos \langle F(\mathbf{r}), F_t \rangle \in \mathbb{R}^N. \quad (5)$$

Although using dense pixel-level CLIP feature improves efficiency, it also includes a lot of noise. To address this issue, we present a novel adapter  $f_{\theta_a}$  that includes two components, an attention module and a small bottleneck layer. We aim to reconstruct the dense CLIP feature enabling it to reduce noise:

$$F(\mathbf{r})_r \approx f_{\theta_a}(F(\mathbf{r})), \quad (6)$$

we can then use the cosine similarity to optimize  $f_{\theta_a}$ :

$$\mathcal{L}_r = \cos \langle F(\mathbf{r})_r, (F(\mathbf{r})) \rangle. \quad (7)$$

The new dense pixel-level CLIP feature captured is added with the original features via residual connections:

$$F(\mathbf{r})_r^{\text{new}} = \alpha F(\mathbf{r})_r + (1 - \alpha) F(\mathbf{r}). \quad (8)$$

However, this optimization is ineffective because we have not made any fine-tuning of CLIP. We observe that after a certain number of training steps, the rendered CLIP features are difficult to improve due to the limitations of CLIP features. Additionally, the rendered CLIP features are cleaner than CLIP features but sometimes lack some semantic information. Therefore, we propose a self-cross-training strategy that aims to reduce noise and improve rendered CLIP features by optimizing adapters to achieve better performance. To sum up, the  $f_{\theta_a}$  can be found via solving:

$$\begin{aligned} \mathcal{L}_r = & -\beta \left( \frac{f_{\theta_a}(F(\mathbf{r})) \cdot F(\mathbf{r})}{|f_{\theta_a}(F(\mathbf{r}))| |F(\mathbf{r})|} + \frac{f_{\theta_a}(\hat{F}(\mathbf{r})) \cdot \hat{F}(\mathbf{r})}{|f_{\theta_a}(\hat{F}(\mathbf{r}))| |\hat{F}(\mathbf{r})|} \right) \\ & - (1 - \beta) \left( \frac{f_{\theta_a}(\hat{F}(\mathbf{r})) \cdot F(\mathbf{r})}{|f_{\theta_a}(\hat{F}(\mathbf{r}))| |F(\mathbf{r})|} + \frac{f_{\theta_a}(F(\mathbf{r})) \cdot \hat{F}(\mathbf{r})}{|f_{\theta_a}(F(\mathbf{r}))| |\hat{F}(\mathbf{r})|} \right), \end{aligned} \quad (9)$$

where  $\beta$  is a weight, and we find through experiments that setting it to 0.3 is optimal. The noise can be mitigated by optimizing adapters and the rendered CLIP feature is cleaner than the original CLIP feature, but its semantic information at the edges is blurred.

### 3.2 Low-Rank Transient Query Attention

For the problem of blurred edges of rendered CLIP features, an intuitive idea is to utilize convolution or attention mechanisms, applying them directly to rendered CLIP features may result in a loss of spatial information [5]. Hence, we try to enhance these features by finding relationships between 3D points. There are some challenges in using convolution or attention mechanisms directly on 3D points. First of all, due to possible occlusion problems in space, the 3D points participating in the training are not fixed at each step. Secondly, the large number of 3D points leads to a significant increase in complexity, which is unacceptable. To address this challenge, we consider the use of low-rank decomposition to streamline the computation within attention mechanisms [56]. By reducing the rank of the matrix, this approach effectively reduces computing complexity and storage requirements, enabling representation through a reduced number of parameters. In attention mechanisms, this allows for an approximation of the original attention matrix by a lower-dimensional equivalent, preserving essential information with minimal loss. The primary benefit of low-rank decomposition lies in its efficiency in cutting down computational demands, a crucial factor when dealing with extensive data sets. By breaking down the initial high-dimensional attention matrix into two smaller matrices, we effectively lessen both the size and computational load of the model, whilst retaining its critical features.

Given a batch of sampling points, we can obtain their corresponding features from the feature grid, represented as

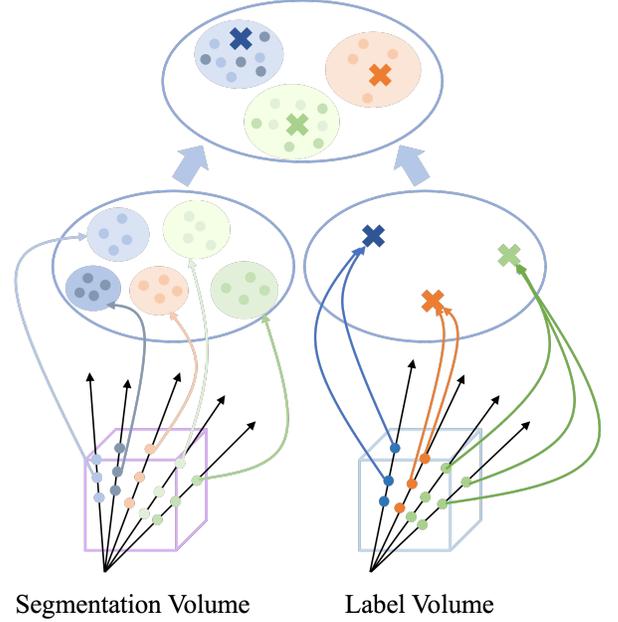


Fig. 4. **Employing label volume to generate cluster centroids.** Our method progressively aggregates points lying on the same ray into a shared cluster centroid during training. This process effectively groups 3D points, which are spatially represented on similar-looking features, into the same category. As a result, 3D points that share close appearances in their feature are associated with the same cluster, reinforcing the consistency of the categorization based on their color similarities.

$F_{sp} \in \mathbb{R}^{S \times D}$ ,  $S$  is the number of sampling points in a batch (i.e., the product of the number of rays and the number of sampling points along each ray), and  $D$  is the number of feature channels. For the vanilla self-attention mechanism:

$$Q = q(F_{sp}), K = k(F_{sp}), V = v(F_{sp}), \quad (10)$$

where  $q, k, v$  represent  $1 \times 1$  convolution layers. We can then get the  $\hat{F}_{sp}$ :

$$\hat{F}_{sp} = V \otimes \text{Softmax}(Q \otimes K). \quad (11)$$

The complexity can be expressed as  $\mathcal{O}(S^2D)$ , where  $S$  is a large number, typically exceeding 20,000. Consequently, employing vanilla attention with a global receptive field results in overwhelming computation complexity.

For a matrix  $W \in \mathbb{R}^{d \times k}$ , we represent it by a low-rank decomposition  $W = BA$ , where  $B \in \mathbb{R}^{d \times r}$ ,  $A \in \mathbb{R}^{r \times k}$ , and the rank  $r \ll \min(d, k)$ . This decomposition allows us to approximate the attention matrix efficiently with reduced dimensions. Thus, we propose a novel low-rank transient query attention, leveraging a transient query learnable parameter  $T \in \mathbb{R}^{s \times D}$ . For vanilla attention output  $\hat{F}_{sp}$ , we can regard it as  $W_F \in \mathbb{R}^{S \times D}$  and factorized it:

$$\begin{aligned} W_F &= \underbrace{M_{qk} \circ v_v}_{\text{Vanilla Attention}} \\ &= \underbrace{M_{qT} \circ M_{kT} \circ v_v}_{\text{Transient Query Attention}} \end{aligned} \quad (12)$$

where  $M_{qk} \in \mathbb{R}^{S \times S}$ ,  $M_{qT} \in \mathbb{R}^{S \times s}$ ,  $M_{kT} \in \mathbb{R}^{s \times S}$ , and  $v_v \in \mathbb{R}^{S \times D}$ . The  $M_{qT}$  and  $M_{kT}$  can approximate the  $M_{qk}$ , and

TABLE 1

**Quantitative comparison** across competing methods on the 3D-OVS [6], where overall indicates the average results across all the metrics in the three different scenes. We highlight the **best**, **second-best**, and **third-best** scores.

Methods	Training Time	bed		sofa		lawn		room		bench		table		overall	
		mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.
LSeg [45]	-	56.0	87.6	04.5	16.5	17.5	77.5	19.2	46.1	06.0	42.7	07.6	29.9	18.5	50.1
ODISE [49]	-	52.6	86.5	48.3	35.4	39.8	82.5	52.5	59.7	24.1	39.0	39.7	34.5	42.8	56.3
OV-Seg [57]	-	79.8	40.4	66.1	69.6	81.2	92.1	71.4	49.1	88.9	89.2	80.6	65.3	78.0	67.6
DFE [7]	184 min	56.6	86.9	03.7	09.5	42.9	82.6	25.1	51.4	06.1	42.8	07.9	30.1	23.7	50.6
LERF [5]	54 min	73.5	86.9	27.0	43.8	73.7	93.5	46.6	79.8	53.2	79.7	33.4	41.0	51.2	70.8
3D-OVS [6]	158 min	89.5	96.7	74.0	91.6	88.2	97.3	92.8	98.9	89.3	96.3	88.8	96.5	87.1	96.2
Our	11 min	91.4	98.0	86.0	96.2	88.5	98.4	85.9	97.7	88.3	96.9	88.5	96.3	88.1	97.3

$s \ll \min(S, D)$ . Specifically, we can use  $T$  to query  $Q$  and  $K$  respectively and calculate Softmax:

$$\hat{T}_q = \text{Softmax}(T \otimes Q) \in \mathbb{R}^{s \times s}, \quad (13)$$

$$\hat{T}_k = \text{Softmax}(T \otimes K) \in \mathbb{R}^{s \times s}. \quad (14)$$

Next, we apply  $\hat{T}_k$  to query  $V$  and calculate Softmax as the  $V$  of vanilla attention:

$$\hat{V} = \text{Softmax}(\hat{T}_k \otimes V) \in \mathbb{R}^{s \times D}. \quad (15)$$

We then can get the  $\hat{F}_{sp}$ :

$$\hat{F}_{sp} = \hat{T}_q \otimes \hat{V} \in \mathbb{R}^{s \times D}. \quad (16)$$

Consequently, the complexity is  $\mathcal{O}(SsD)$ . This reduction in complexity significantly contributes to a more efficient training process, thereby effectively minimizing training time.

### 3.3 Label Volume for Enhance Consistency

In 3D scenes, maintaining consistent views is crucial. To this end, we propose a novel label volume to transform the segmentation problem into a classification problem. Specifically, each 3D point  $\mathbf{x} \in \mathbb{R}^3$  yields a label vector  $L_{\mathbf{x}} \in \mathbb{R}^3$  from label volume  $L$ . Then we can render the label vector of each ray  $\mathbf{r}$  using volume rendering:

$$L(\mathbf{r}) = \text{Softmax}\left(\sum_{i=1}^N w_i L_i\right) \in [0, 1]^D. \quad (17)$$

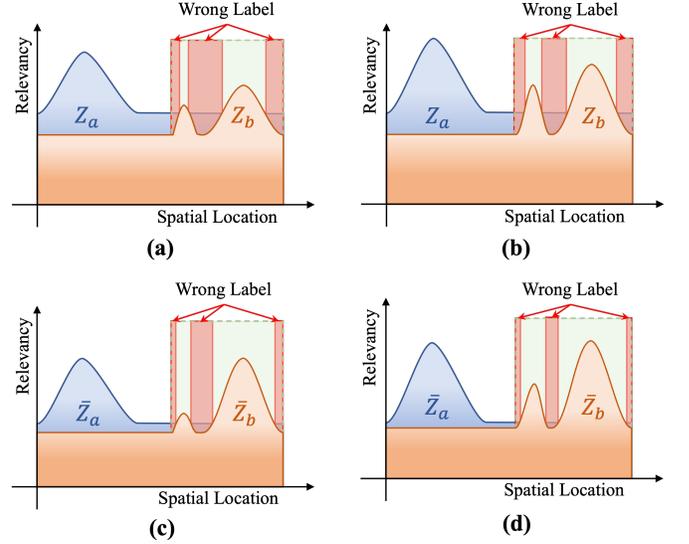
where  $w_i$  refers to Equation (2), we employ a Softmax function on the label vector for each ray, ensuring that the probabilities sum up to 1. We can calculate pseudo labels  $L_{pGT}$  using classes' text features  $F_t$ :

$$L_{pGT} = \text{cos}\langle L(\mathbf{r}), F_t \rangle \in \mathbb{R}^N. \quad (18)$$

We can then use cross-entropy to optimize  $z_I(\mathbf{r})$  and  $z(\mathbf{r})$ :

$$\mathcal{L}_{CE} = - \sum_{i=1}^N L_{pGT_i} \cdot (\log(z_I(\mathbf{r})_i) + \log(z(\mathbf{r})_i)), \quad (19)$$

where  $N$  represents the number of classes. As illustrated in Section 3.1,  $z_I(\mathbf{r})$  include certain noise and  $z(\mathbf{r})$  may lose some semantic information. However,  $L_{pGT}$  is a learnable smoothing pseudo label, during the backpropagation process, these noisy and semantically inaccurate areas can contribute to incorrect gradients, leading to inaccuracies in labels, as illustrated in Figure 4. To alleviate this issue, we adopt the concept of ensemble learning, weighting  $z_I(\mathbf{r})$  and



**Fig. 5. Mitigating the ambiguity in CLIP features.** We employ a simplified text augmentation strategy to standardize relevance maps. Observing the original relevance maps  $Z_a$  and  $Z_b$  in (a), we note that the relevance of class **a** within the red-highlighted area is higher than in other image regions. Due to the higher absolute relevance of class **a** in this area, the ambiguity of CLIP features results in the red region being classified as class **a**, even though class **b** is also present. In (b), we reduce this ambiguity by simply repeating the text to recalculate the relevance maps  $Z_a$  and  $Z_b$ , thereby enhancing the accuracy of regional class assignments. In (c), standardizing the relevance maps of each class to a fixed range also can reduce ambiguity. In (d), we combine text repetition with standardization of the relevance maps, significantly reducing the ambiguity in classification and leading to more precise regional class allocations.

$z(\mathbf{r})$ , and optimizing the smoothed pseudo-label  $L_{pGT}$ . The expression for Equation (19) can be expressed as follows:

$$\mathcal{L}_{CE} = - \sum_{i=1}^N L_{pGT_i} \cdot (\log(z_I(\mathbf{r})_i) + \log(z(\mathbf{r})_i) + \log(\gamma z_I(\mathbf{r})_i + (1 - \gamma)z(\mathbf{r})_i)), \quad (20)$$

where  $\gamma$  control  $z_I(\mathbf{r})$  and  $z(\mathbf{r})$  weight. By minimizing  $L_{sGT}$ , we ensure that 3D points on each ray in the label volume are aggregated onto their most probable category, causing the 3D point set on the segmentation branch to concentrate on relatively correct classifications. Relatively accurate classification of a given 3D point ensures consistency across views during rendering. Although converting the classification problem can obtain relative accuracy segmentation, it



Fig. 6. **Qualitative comparison.** Visualization of segmentation results in 3D-OVS dataset 3 scenes [6]. Note that we generate object masks in this dataset requiring the complete category list.

cannot handle the ambiguity between original CLIP features and textual correlations, as illustrated in Figure 5.

### 3.4 Simplified Text Augmentation for Mitigating Ambiguity

To alleviate the ambiguity between original CLIP features and textual correlations, we propose a simplified text augmentation strategy. Specifically, we let the text descriptions  $\mathcal{T}$  repeat twice to obtain  $\mathcal{T}_{aug} = \{t_i \times 2\}_{i=1}^N$ , and then get the classes' text features  $F_t^{aug} = E_t(\mathcal{T}_{aug}) \in \mathbb{R}^{N \times D}$ . Given an image, we can get its dense pixel-level CLIP feature  $F_I \in \mathbb{R}^{H \times W \times D}$  from the CLIP image encoder and though the adapter. We can then get the image segmentation logits  $z_I$ :

$$z_I = \cos \langle F_I, F_t^{aug} \rangle \in \mathbb{R}^{N \times H \times W}. \quad (21)$$

We independently normalize the correlation of each class in the input view to  $[0, 1]$  to alleviate the ambiguity of CLIP features:

$$\bar{z}_I = \frac{z_I - \min(z_I)}{\max(z_I) - \min(z_I)} \in [0, 1]^{N \times H \times W}. \quad (22)$$

Here, the functions  $\min(\cdot)$  and  $\max(\cdot)$  are utilized to retrieve the minimum and maximum values within the spatial dimensions (i.e.  $H$  and  $W$ ). We can then further optimize the segmentation logits  $z_I(\mathbf{r})$  of the ray  $\mathbf{r}$ :

$$\mathcal{L}_{aug} = \cos \langle z_I(\mathbf{r}), \bar{z}_I(\mathbf{r}) \rangle. \quad (23)$$

Note that  $\bar{z}_I(\mathbf{r})$  is to select the image pixels corresponding to the batch of rays  $\mathbf{r}$ .

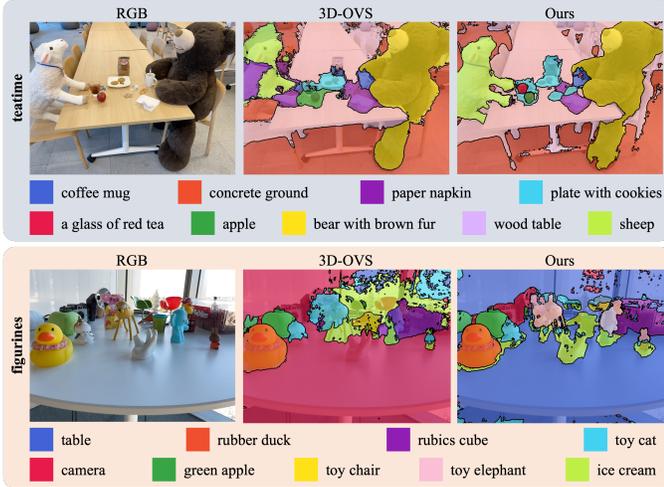


Fig. 7. **Qualitative comparison.** Visualization of segmentation results in LERF [5] dataset.

## 4 EXPERIMENT

### 4.1 Implementation Details

To reconstruct a 3D scene from multiview images, we follow the same training settings as TensoRF [36] and 3D-OVS [6]. During the segmentation phase of training, we train the model for 15,000 iterations. For the first 5,000 iterations, we freeze the parameters of the appearance volume and density volume and train only on the label volume, CLIP feature branches, and adapters. In the subsequent 10,000 iterations, we further fine-tune the appearance volume and RGB branches. At the 8,000th iteration, we introduce a self-crossing loss function to optimize the adapter. We used the Adam optimizer with  $\beta_{tas} = (0.9, 0.99)$ . The initial learning rates used to train the volumetric and segmentation MLP branches are 0.02 and  $1e-4$  respectively. During the fine-tuning phase, these learning rates are adjusted to  $5e-3$  and  $5e-5$ . In addition, we also implemented a learning rate decay strategy with a decay coefficient of 0.1. The model dynamically computes dense pixel-level CLIP features for each training view during training. We use the ViT-B/16 CLIP model to extract the image and text features. The model runs on an NVIDIA 3090 GPU with 24G of memory. The average training time is approximately 11 minutes, and the memory required during training does not exceed 10G.

### 4.2 Evaluation Metrics

In this work, we adopt two metrics, the mIoU score and the accuracy score, to evaluate the performance of our image segmentation model. mIoU score is a standard indicator to measure the segmentation effect of the model. It calculates the average of the ratio of intersection and union between the segmentation area predicted by the model and the true segmentation area. The accuracy score measures the proportion of pixels correctly classified by the model, that is, the number of pixels correctly predicted by the model as a proportion of the total number of pixels.

### 4.3 Dataset

To assess the efficacy of our approach, we primarily focus on the 3D-OVS [6] dataset, which is specially designed for

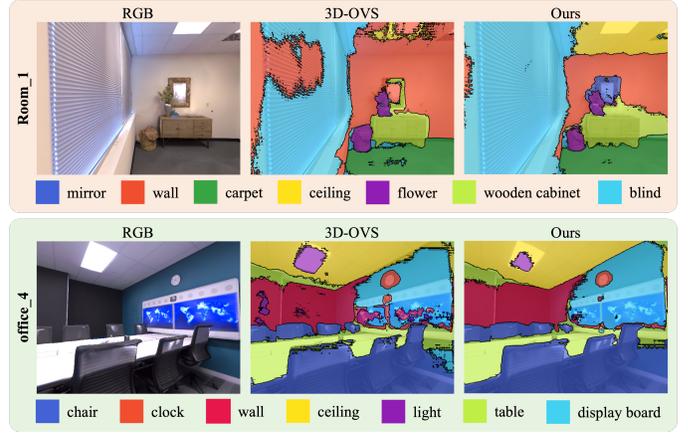


Fig. 8. **Qualitative comparison.** Visualization of segmentation results in Replica [58] dataset.

TABLE 2  
Comparison of inference time and FPS for our and 3D-OVS.

Method	Inference Time (s/Frame)	FPS
3D-OVS	6.156	0.16
Our	9.125	0.11

open-vocabulary 3D semantic segmentation and features a diverse collection of long-tail objects in various poses and backgrounds, complete with a detailed list of categories. In addition to the 3D-OVS dataset, we further validate our method using two supplementary datasets. The first is the widely-used Replica [58] dataset, targeting multi-view indoor scenes. The second is the LERF [5] dataset, which is derived from complex real-world scenes captured via the iPhone App Polycam.

### 4.4 Inference Cost

We compared the inference time and FPS between our method and 3D-OVS [6]. As shown in Table 2, due to the retention of the adapter and low-rank transient query attention module in the inference phase, our inference time and FPS have decreased. However, it should be emphasized that the main focus of this work is to reduce the consumption of the training phase rather than the consumption of the inference phase.

### 4.5 Comparison with State-of-the-Art Methods

#### 4.5.1 Qualitative Results

In Figures Figure 6, Figure 7, and Figure 8, we present the visualization results of three datasets. For the 3D-OVS [6] dataset, we demonstrate the segmentation effects observed from different perspectives, as shown in Figure 6. Compared to other 3D methods, our methods demonstrate superior performance in capturing fine edge details. For the LERF [5] and Replica [58] datasets, we compared our method with the current state-of-the-art 3D open vocabulary segmentation method, 3D-OVS. In terms of processing complex scenes, our method can more precisely identify scene details, such as the apples and teacups in the teatime scene, and the mirror in the room\_1 scene. These qualitative results prove the

TABLE 3  
Quantitative comparison across competing methods on the LERF [5] and Replica [58]. We highlight the **best** scores.

Methods	LERF [5]								Replica [58]							
	<i>teatime</i>		<i>figurines</i>		<i>ramen</i>		<i>waldo_kitchen</i>		<i>room_0</i>		<i>room_1</i>		<i>office_3</i>		<i>office_4</i>	
	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.	mIoU	Acc.
LERF [5]	26.0	82.8	34.6	55.0	25.2	68.3	33.9	42.7	7.7	38.2	12.8	29.7	12.6	37.2	8.2	16.3
3D-OVS [6]	56.1	87.7	44.8	77.3	28.7	70.2	39.3	45.6	12.5	40.1	13.1	30.2	12.1	36.1	15.8	28.9
Our	62.4	88.0	63.5	83.9	44.6	75.8	41.3	48.9	24.3	54.7	25.0	45.9	18.7	64.3	39.7	62.5

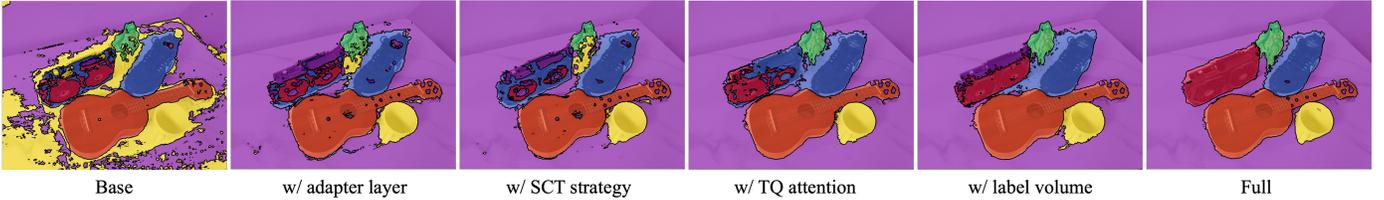


Fig. 9. Ablation studies visualization results on the table scene of 3D-OVS [6]. We present the visualization results of the adapter layer, the label volume, the self-cross-training (SCT) strategy, the low-rank transient query (TQ) attention, and the full module.

advantages of our proposed method in generating accurate object boundaries and ensuring 3D semantic consistency.

#### 4.5.2 Quantitative Results

In Table 1, we benchmark our method using three 2D open vocabulary segmentation methods [45], [49], [57] and three NeRF-based methods capable of 3D open vocabulary segmentation: DFF [7], LERF [5], and 3D-OVS [6]. DFF is a pioneering work in the field of 3D open vocabulary segmentation, which applies the feature maps of the 2D open vocabulary segmentation method LSeg [45] to achieve its segmentation results. LERF, 3D-OVS are closely aligned with our proposed methods as they use knowledge distillation from CLIP and DINO. We observe that our method significantly surpasses existing 2D and 3D methods in multiple aspects. Our method not only outperforms 2D-based methods such as ODISE [49] and OV-Seg [57] in performance, but also outperforms 3D-based methods including LERF [5] and 3D-OVS [6]. Specifically, our method performs better than the current state-of-the-art 3D-OVS in 3 out of 6 scenes in the 3D-OVS dataset, and is comparable to 3D-OVS in 1 scene. In particular, significant performance improvements are achieved in the sofa scene. Although 3D-OVS adopts the optimization strategy of multi-scale pixel-level CLIP features and DINO features, the overall performance of our method is outstanding in terms of average mIoU score and accuracy. In addition, from an efficiency perspective, 3D-OVS consumes a lot of resources during the training phase, and the optimization of a single scene takes 158 minutes. In comparison, our method only takes 11 minutes for single-scene optimization, and is 14.3 times more efficient than 3D-OVS. In Table 3, we compare our method against LERF [5] and 3D-OVS [6] using the LERF [5] and Replica [58] datasets. Our method consistently demonstrates superior performance across the 8 scenarios evaluated. This quantitative experiment demonstrates our method’s robust capability in executing precise NeRF-based 3D word segmentation, relying on a streamlined architecture.

## 4.6 Analysis

In this subsection, we performed an ablation study for different components of our Laser framework. In Table 4, we initially start distilling dense CLIP features to the segmentation branch as the base. Subsequently, we gradually introduce other components, demonstrating the changes in mIoU scores and accuracy scores with the addition of each component based on this foundation. Additionally, we also show the qualitative results of this change (the table scene of 3D-OVS), as shown in the Figure 9.

### 4.6.1 Investigation of Optimizing Adapter Layer

We studied the importance of the designed components in our Laser framework. Specifically, we define the following methods: A) *w/ adapter layer* to directly optimize it use Equation (7), and B) *w/ SCT strategy* to directly evaluate our self-cross training strategy (Equation (9)). The first three columns of Figure 9 show the comparison of the adapter layer, self-cross training strategy with the baseline. From the results, although the segmentation effect can be achieved by directly extracting pixel-level features from distilled dense CLIP, since CLIP itself only has classification functions and lacks segmentation capabilities, a large amount of noise will be generated if the segmentation task is not specifically processed, refer to Section 3. *Method A* can effectively reduce noise, thanks to the features reconstructed through the bottleneck layer, resulting in cleaner output features. *Method B* can further reduce noise. The effectiveness of this strategy is based on our observation that rendered features demonstrate higher clarity compared to reconstructed features. Therefore, by enabling mutual supervision between these two types of features, we facilitate their mutual enhancement. Furthermore, we delved into the significance of the adapter residual connection ratio  $\alpha$  and the ratio  $\beta$  in the self-cross training strategy on the importance to the model. In Table 5, with  $\alpha$  set to 0.2, we observed that the model achieved its optimal performance. An  $\alpha$  value of 0 means the model directly utilizes the original dense CLIP features, while a setting of 1 implies complete reliance on reconstructed features. The results show that when relying

TABLE 4

**Ablation studies** on 3D-OVS [6]. We evaluate the adapter layer, the label volume (LV), the self-cross-training (SCT) strategy, the low-rank transient query (TQ) attention, and the simplified text augmentation loss  $\mathcal{L}_{aug}$ . Our final method in the last row performs the best.

Method	mIoU	Acc.
Base (only distill CLIP feature)	59.4	74.5
+ adapter layer	68.2	83.4
+ adapter layer + SCT	73.9	89.7
+ adapter layer + SCT + TQ	78.6	92.4
+ adapter layer + SCT + TQ + LV	85.6	96.0
+ adapter layer + SCT + TQ + LV + $\mathcal{L}_{aug}$	88.1	97.3

TABLE 5

**Ablations** on varying the residual ratio  $\alpha$  for adapter layer.

Ratio $\alpha$	0	0.2	0.4	0.6	0.8	1.0
mIoU	81.8	88.1	87.6	87.3	87.2	71.8
Acc.	94.3	97.3	96.2	96.0	95.9	90.6

entirely on reconstructed features, there is a significant decrease in model performance. This indicates that although reconstructed features help reduce the impact of noise, they also result in the loss of multimodal information in the CLIP features. While optimizing adapters through the method described in Equation (7) can mitigate the impact of noise to a certain extent, the supervision signal it relies on (that is, the original dense CLIP features) is inherently filled with a substantial amount of noise. In contrast, the rendered features are clearer, providing more effective supervision. Therefore, we attempt to use a self-cross training strategy as detailed in Equation (9) to optimize the adapter, aiming to achieve better performance. In Table 6, we explore the impact of different  $\beta$  on model performance. Among them, when  $\beta = 0.3$ , the model has the best performance and has a greater improvement compared to other hyperparameters.

#### 4.6.2 Investigation of Low-Rank Transient Query Attention

To investigate the complexity and performance of our low-rank transient query attention, we compared vanilla self-attention and without any attention block. As shown in the fourth column of Figure 9, our low-rank transient query attention can enhance the information details of object edges, as we discussed in Section 3. However, for categories with similar colors, this approach may lead to additional misclassifications. For example, the black edges of graphics cards and black shoes are misidentified as the same category. The results are shown in Table 7, where D represents the dimensionality of the attention projection. By increasing the dimension of the hidden layer from 16 to 64, we found that excessively small dimensions significantly reduce model performance. The reason is that lower dimensions impair the capability to express semantic information of features.

TABLE 6

**Ablations** on varying the ratio  $\beta$  for  $\mathcal{L}_r$ .

Ratio $\beta$	0	0.1	0.3	0.5	0.7	0.9	1.0
mIoU	87.2	87.7	88.1	87.8	86.5	85.9	85.8
Acc.	95.9	96.8	97.3	96.3	95.7	95.1	95.0

TABLE 7

**Ablations** on varying the low-rank transient query (TQ) attention, and Vanilla is the Vanilla self-attention.

Method	mIoU	Acc.	FLOPs ↓	# Params ↓	Train Time ↓
W/o attention	85.31	95.41	-	-	10 min
Vanilla (D=32)	87.98	97.13	26.91G	0.25M	2.2 hrs
TQ (D=16)	37.18	83.58	0.69G	0.13M	11 min
TQ (D=32)	88.10	97.25	1.48G	0.26M	11 min
TQ (D=48)	86.33	95.27	2.33G	0.39M	11 min
TQ (D=64)	86.32	95.55	3.28G	0.53M	11 min

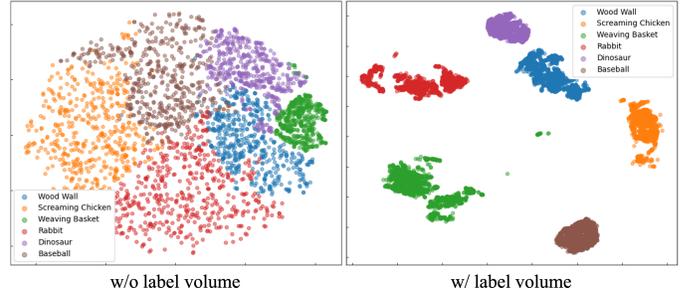


Fig. 10. **Ablation** on the label volume. 3D points with similar feature expressions are gathered into the same category in the feature space, thereby enhancing the classification effect based on light consistency.

However, increasing the dimension also leads to a decline in performance, as excessively high dimensions introduce unnecessary redundancy. The optimal bottleneck dimension is 32, which can retain sufficient semantics without redundancy. The complexity of these two types of attention refers to Section 3. From the results, under the condition of maintaining the same dimensions, the traditional vanilla self-attention model has 18.18 times the floating-point operations (FLOPs) compared to our proposed low-rank transient query attention model, and 12 times our model in training time. Although our model has only 0.01MB more parameters than the vanilla self-attention, our mIoU score is improved by 0.02%, and the accuracy score is improved by 0.12%. Furthermore, compared to the baseline model without any integrated attention block, our model only adds 1 minute to the training time, while significantly improving the mIoU score and accuracy.

#### 4.6.3 Investigation of Label Volume

We utilized label volumes to transform the segmentation task into a classification problem, a change that is crucial for enhancing model performance. As shown in the fifth column of Figure 9, the introduction of label volumes significantly optimized the segmentation results. In Figure 10, we found that under the influence of label volumes, features of different categories exhibited a more compact and distinct distribution. This indicates that label volumes effectively improved the model's ability to discriminate the similarity

TABLE 8

**Ablations** on varying the ratio  $\gamma$  for  $\mathcal{L}_{CE}$ .

Ratio $\gamma$	0.1	0.3	0.5	0.7	0.9
mIoU	84.1	87.3	88.1	86.5	81.6
Acc.	94.4	95.9	97.3	95.7	93.1

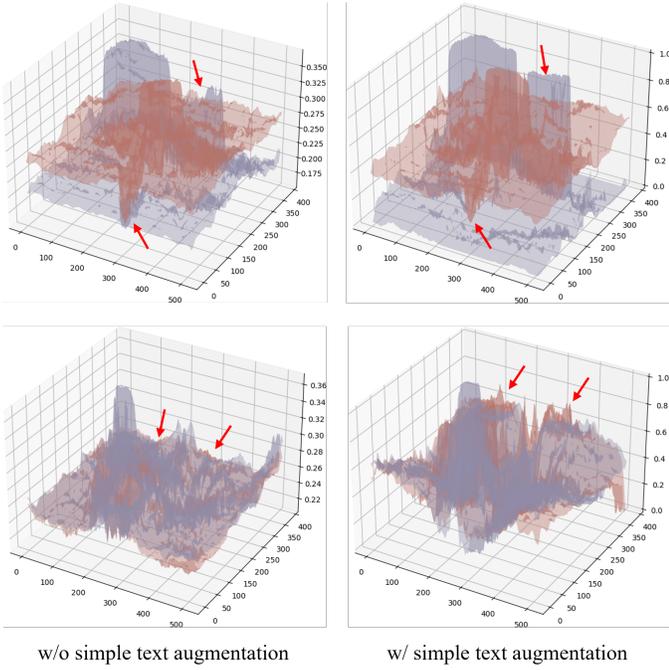


Fig. 11. **Ablation** on simplified text augmentation. The x-axis and y-axis represent the pixel position of the image, while the z-axis indicates the degree of similarity between the text and the image. *w/o simplified text augmentation*, there is notable ambiguity within the CLIP features indicated by the red arrows. Conversely, *w/ simplified text augmentation*, the ambiguity in the CLIP features in these areas is reduced.

between features of the same category, thereby making the feature representation of the same category more unified. This enhanced unity has led to a significant improvement in model performance. Furthermore, we delved into the specific impact of label volume at different weights on model performance. As shown in Table 8, we observed that the model achieved optimal performance when  $\gamma$  is set to 0.5. This can be attributed to the fact that label volume is achieved through the joint optimization of the original CLIP features and the rendered features. Overreliance on the original CLIP features may introduce excessive noise; conversely, overly relying on the rendered features may weaken the model’s multimodal capability, thereby affecting its segmentation efficacy.

#### 4.6.4 Investigation of Simplified Text Augmentation Strategy

As shown in the last column of Figure 9, after adopting a simplified text augmentation strategy, the errors inside the graphics cards are significantly reduced, and the segmentation effect on the edges of the graphics cards is also noticeably improved. This is because the text augmentation strategy enhances the distinction between categories through repeated text enhancements and, after normalization, effectively alleviates the ambiguity problem caused by the high similarity between categories. The results of Figure 11 further confirm the effectiveness of our method: *w/o simplified text augmentation strategy*, the locations marked with red arrows had extremely high category similarity, making it difficult to distinguish. However, *w/ simplified*

*text augmentation strategy*, the categories within the areas indicated by the red arrows are identified and differentiated.

## 4.7 Limitations

While our approach demonstrates notable enhancements in terms of speed and storage resource utilization when compared to prior methodologies, it is largely dependent on the predictive capabilities of dense pixel-level CLIP characteristics. The performance of our model is significantly influenced by the quality of these dense features, which frequently exhibit noise. As evidenced by our qualitative study, despite implementing many strategies to mitigate noise, our findings remain less refined and exhibit a certain level of noise in comparison to previously employed multi-scale features. This limitation becomes particularly evident when dealing with small objects or objects of similar colors. This suggests that while our approaches have made strides in noise reduction, the precision in identifying and segmenting objects under these challenging conditions requires further enhancement. Additionally, the reason for the need to modify CLIP’s encoder is due to the requirement of predicting dense pixel-level CLIP characteristics without any fine-tuning. While this technique retains its multi-modal capabilities, it is important to acknowledge that these adjustments will obviously have an impact on the performance of the model. Furthermore, our approach is limited in its ability to enhance model performance due to its inability to fine-tune the encoder utilizing extensive data sets, such as the 2D segmentation task.

## 5 CONCLUSION

In this work, we propose Laser, an efficient language-guided segmentation of 3D radiation fields. Compared with previous research, we optimized the workflow and used dense pixel-level CLIP feature distillation to guide the radiation field to have text-guided segmentation capabilities. Specifically, we first combine adapter and self-cross-training strategies to effectively mitigate the noise introduced by dense pixel-level CLIP features. Next, we introduced low-rank transient query attention, which not only effectively enhanced the edge clarity of the segmentation map, but also reduced the complexity of vanilla self-attention in 3D feature processing. Furthermore, we enhance the consistency across different views by transforming the segmentation problem into a classification problem. Finally, we adopt a simplified text enhancement strategy to alleviate the ambiguity of CLIP correlations. Extensive experiments demonstrate that our method outperforms current state-of-the-art technologies in both speed and performance.

Our method is based on TensorRF [36] implementation and is limited in rendering speed, but a potential solution is to use state-of-the-art 3D Gaussian splatting technology [2]. Furthermore, as we discussed in the limitation analysis, distilling directly from dense pixel-level CLIP features with noise will inevitably affect model performance. Therefore, a focus of our future research is to explore how to denoise these features more effectively. On the other hand, we plan to explore more general graph-based multimodal methods to more comprehensively capture the global relationship

between different modalities. Currently, our method mainly focuses on the mining of local modal relationships but lacks the overall modal association. By introducing a more general graph-based method [59], we hope to make up for this deficiency and thus improve the performance of the model in multimodal tasks.

## ACKNOWLEDGMENTS

This work is supported by the UK Medical Research Council (MRC) Innovation Fellowship under Grant MR/S003916/2, International Exchanges 2022 IEC\NSFC\223523.

## REFERENCES

- [1] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [2] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," *ACM Transactions on Graphics*, vol. 42, no. 4, July 2023. [Online]. Available: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- [3] Y. Siddiqui, L. Porzi, S. R. Bulò, N. Müller, M. Nießner, A. Dai, and P. Kotschieder, "Panoptic lifting for 3d scene understanding with neural fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9043–9052.
- [4] Y. Bhalgat, I. Laina, J. F. Henriques, A. Zisserman, and A. Vedaldi, "Contrastive lift: 3d object instance segmentation by slow-fast contrastive fusion," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [Online]. Available: <https://openreview.net/forum?id=b4bbbov4Xu>
- [5] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik, "Lerf: Language embedded radiance fields," in *International Conference on Computer Vision (ICCV)*, 2023.
- [6] K. Liu, F. Zhan, J. Zhang, M. Xu, Y. Yu, A. El Saddik, C. Theobalt, E. Xing, and S. Lu, "Weakly supervised 3d open-vocabulary segmentation," in *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [7] S. Kobayashi, E. Matsumoto, and V. Sitzmann, "Decomposing nerf for editing via feature field distillation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 23 311–23 330, 2022.
- [8] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [9] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [10] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.
- [11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [12] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021.
- [13] Z. T. Zheng Ding, Jieke Wang, "Open-vocabulary universal image segmentation with maskclip," in *International Conference on Machine Learning*, 2023.
- [14] S. Cho, H. Shin, S. Hong, S. An, S. Lee, A. Arnab, P. H. Seo, and S. Kim, "Cat-seg: Cost aggregation for open-vocabulary semantic segmentation," 2023.
- [15] F. Wang, J. Mei, and A. Yuille, "Sclip: Rethinking self-attention for dense vision-language inference," *arXiv preprint arXiv:2312.01597*, 2023.
- [16] J.-X. Chai, X. Tong, S.-C. Chan, and H.-Y. Shum, "Plenoptic sampling," in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, 2000, pp. 307–318.
- [17] N. K. Kalantari, T.-C. Wang, and R. Ramamoorthi, "Learning-based view synthesis for light field cameras," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, pp. 1–10, 2016.
- [18] S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 43–54.
- [19] M. Levoy and P. Hanrahan, "Light field rendering," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 31–42.
- [20] G. Wu, B. Masia, A. Jarabo, Y. Zhang, L. Wang, Q. Dai, T. Chai, and Y. Liu, "Light field image processing: An overview," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 7, pp. 926–954, 2017.
- [21] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen, "Unstructured lumigraph rendering," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001, pp. 425–432.
- [22] P. E. Debevec, C. J. Taylor, and J. Malik, "Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach," in *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, 1996, pp. 11–20.
- [23] J. Flynn, I. Neulander, J. Philbin, and N. Snavely, "Deepstereo: Learning to predict new views from the world's imagery," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5515–5524.
- [24] P. Hedman, J. Philip, T. Price, J.-M. Frahm, G. Drettakis, and G. Brostow, "Deep blending for free-viewpoint image-based rendering," *ACM Transactions on Graphics (TOG)*, vol. 37, no. 6, pp. 1–15, 2018.
- [25] P. Hedman, T. Ritschel, G. Drettakis, and G. Brostow, "Scalable inside-out image-based rendering," *ACM Transactions on Graphics (TOG)*, vol. 35, no. 6, pp. 1–11, 2016.
- [26] J. Kopf, M. F. Cohen, and R. Szeliski, "First-person hyper-lapse videos," *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, pp. 1–10, 2014.
- [27] E. Penner and L. Zhang, "Soft 3d reconstruction for view synthesis," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 6, pp. 1–11, 2017.
- [28] G. Riegler and V. Koltun, "Free view synthesis," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIX 16*. Springer, 2020, pp. 623–640.
- [29] —, "Stable view synthesis," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 12 216–12 225.
- [30] L. Liu, J. Gu, K. Zaw Lin, T.-S. Chua, and C. Theobalt, "Neural sparse voxel fields," *Advances in Neural Information Processing Systems*, vol. 33, pp. 15 651–15 663, 2020.
- [31] B. Mildenhall, P. Hedman, R. Martin-Brualla, P. P. Srinivasan, and J. T. Barron, "Nerf in the dark: High dynamic range view synthesis from noisy raw images," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16 190–16 199.
- [32] W. Xian, J.-B. Huang, J. Kopf, and C. Kim, "Space-time neural irradiance fields for free-viewpoint video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 9421–9431.
- [33] Y. Xiangli, L. Xu, X. Pan, N. Zhao, A. Rao, C. Theobalt, B. Dai, and D. Lin, "Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering," in *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*. Springer, 2022, pp. 106–122.
- [34] Q. Xu, Z. Xu, J. Philip, S. Bi, Z. Shu, K. Sunkavalli, and U. Neumann, "Point-nerf: Point-based neural radiance fields," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5438–5448.
- [35] A. Yu, V. Ye, M. Tancik, and A. Kanazawa, "pixelnerf: Neural radiance fields from one or few images," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4578–4587.
- [36] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su, "Tensorf: Tensorial radiance fields," in *European Conference on Computer Vision*. Springer, 2022, pp. 333–350.
- [37] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," *ACM Transactions on Graphics (ToG)*, vol. 41, no. 4, pp. 1–15, 2022.
- [38] C. Sun, M. Sun, and H.-T. Chen, "Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction," in *Pro-*

- ceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5459–5469.
- [39] A. Yu, R. Li, M. Tancik, H. Li, R. Ng, and A. Kanazawa, “Plenotrees for real-time rendering of neural radiance fields,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5752–5761.
- [40] S. Fridovich-Keil, A. Yu, M. Tancik, Q. Chen, B. Recht, and A. Kanazawa, “Plenoxels: Radiance fields without neural networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 5501–5510.
- [41] M. Bucher, T.-H. Vu, M. Cord, and P. Pérez, “Zero-shot semantic segmentation,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [42] Y. Xian, S. Choudhury, Y. He, B. Schiele, and Z. Akata, “Semantic projection network for zero-and few-label semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8256–8265.
- [43] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” *arXiv preprint arXiv:1301.3781*, 2013.
- [44] J. Xu, S. De Mello, S. Liu, W. Byeon, T. Breuel, J. Kautz, and X. Wang, “Groupvit: Semantic segmentation emerges from text supervision,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 18134–18144.
- [45] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, “Language-driven semantic segmentation,” 2022.
- [46] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, “Scaling open-vocabulary image segmentation with image-level labels,” in *European Conference on Computer Vision*. Springer, 2022, pp. 540–557.
- [47] X. Fu, S. Zhang, T. Chen, Y. Lu, L. Zhu, X. Zhou, A. Geiger, and Y. Liao, “Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation,” in *2022 International Conference on 3D Vision (3DV)*. IEEE, 2022, pp. 1–11.
- [48] Q. Shen, X. Yang, and X. Wang, “Anything-3d: Towards single-view anything reconstruction in the wild,” *arXiv preprint arXiv:2304.10261*, 2023.
- [49] J. Xu, S. Liu, A. Vahdat, W. Byeon, X. Wang, and S. De Mello, “Open-vocabulary panoptic segmentation with text-to-image diffusion models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2955–2966.
- [50] S. Zhi, T. Laidlow, S. Leutenegger, and A. J. Davison, “In-place scene labelling and understanding with implicit scene representation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15838–15847.
- [51] V. Tschernezki, I. Laina, D. Larlus, and A. Vedaldi, “Neural feature fusion fields: 3d distillation of self-supervised 2d image representations,” in *2022 International Conference on 3D Vision (3DV)*. IEEE, 2022, pp. 443–453.
- [52] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo *et al.*, “Segment anything,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4015–4026.
- [53] M. Qin, W. Li, J. Zhou, H. Wang, and H. Pfister, “Langsplat: 3d language gaussian splatting,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 20051–20060.
- [54] J.-C. Shi, M. Wang, H.-B. Duan, and S.-H. Guan, “Language embedded 3d gaussians for open-vocabulary scene understanding,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 5333–5343.
- [55] S. Zhou, H. Chang, S. Jiang, Z. Fan, Z. Zhu, D. Xu, P. Chari, S. You, Z. Wang, and A. Kadambi, “Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 21676–21685.
- [56] X. Zhou, C. Yang, H. Zhao, and W. Yu, “Low-rank modeling and its applications in image analysis,” *ACM Computing Surveys (CSUR)*, vol. 47, no. 2, pp. 1–33, 2014.
- [57] F. Liang, B. Wu, X. Dai, K. Li, Y. Zhao, H. Zhang, P. Zhang, P. Vajda, and D. Marculescu, “Open-vocabulary semantic segmentation with mask-adapted clip,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7061–7070.
- [58] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma *et al.*, “The replica dataset: A digital replica of indoor spaces,” *arXiv preprint arXiv:1906.05797*, 2019.

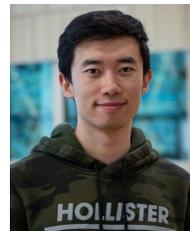
- [59] A. Marcu, M. Pirvu, D. Costea, E. Haller, E. Slusanschi, A. N. Belbachir, R. Sukthankar, and M. Leordeanu, “Self-supervised hypergraphs for learning multiple world interpretations,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 983–992.



**Xingyu Miao** received a master's degree from the School of Information Engineering, Ningxia University, China. Currently, he is a PhD student in the Department of Computing, Durham University, UK. His current main focus is monocular depth estimation and 3D reconstruction.



**Haoran Duan** (Member, IEEE) received a Distinction M.S. degree in Data Science from Newcastle University, UK. He obtained PhD degree in the Department of Computer Science, Durham University. He has been a research associate in Newcastle University, UK, working on deep learning applications. His current research interests focus on the applications/theories of deep learning.



**Yang Bai** received the M.S. degree from Newcastle University, UK, in 2018, and the Ph.D. degree from Durham University, UK, in 2023. He is currently a Research Scientist at the Institute of High Performance Computing (IHPC), ASTAR, Singapore. His current research interests focus on video understanding, multi-modal vision language learning and 3D vision.



**Tejal Shah** received the Ph.D. degree from the School of Computer Science and Engineering, University of New South Wales, Australia. She is a Lecturer in Informatics and member of the Networked and Ubiquitous Systems Engineering (NUSE) group. She's an expert on the research, development, and application of Semantic Web and knowledge representation technologies such as ontologies and vocabularies to address complex information management issues.



**Jun Song** received his Bachelor's and Master's degrees in Computer Science and Technology from China University of Geosciences (Wuhan) in 1997 and 2003, respectively, and his PhD in Information Security from Wuhan University in 2008. He has served as a professor at the School of Computer Science, China University of Geosciences (Wuhan). In 2010, he was a visiting scholar at the University of Victoria, Canada. His research focuses on computer science and information security, with extensive contributions

to these fields.



**Yang Long (Senior Member, IEEE)** is an Associate Professor in the Department of Computer Science, Durham University. He is also an MRC Innovation Fellow aiming to design scalable AI solutions for large-scale healthcare applications. His research background is in the highly interdisciplinary field of Computer Vision and Machine Learning. While he is passionate about unveiling the black-box of AI brain and transferring the knowledge to seek Scalable, Interactable, Interpretable, and sustainable solutions for other disciplinary researches, e.g. physical activity, mental health, design, education, security, and geoengineering. He has authored/coauthored 30+ top-tier papers in refereed journals/conferences such as IEEE TPAMI, TIP, CVPR, AAAI, and ACM MM, and holds a patent and a Chinese National Grant.



**Rajiv Ranjan (Fellow, IEEE)** is a University Chair Professor in Internet of Things research with the School of Computing, Newcastle University, U.K. He is the Director of Networked and Ubiquitous Systems Engineering (NUSE) Group, jointly with Dr. Graham Morgan, School of Computing, Newcastle University. He is also the Academic Director of School of Computing and the Research Director of Newcastle Urban Observatory. He is an internationally established Scientist in the area of distributed systems, having published over 250 scientific papers out of which about 50 papers in IEEE/ACM Transactions journals. He is a fellow of IEEE, Academia Europaea and the Asia-Pacific Artificial Intelligence Association. He is also the Founding Director of the International Centre (U.K. and Australia) on the Internet of Energy (IoE), funded by EPSRC.



**Ling Shao (Fellow, IEEE)** is a Distinguished Professor with the UCAS-Terminus AI Lab, University of Chinese Academy of Sciences, Beijing, China. He was the founder of the Inception Institute of Artificial Intelligence (IIAI) and the Mohamed bin Zayed University of Artificial Intelligence (MBZUAI), Abu Dhabi, UAE. His research interests include generative AI, vision and language, and AI for healthcare. He is a fellow of the IEEE, the IAPR, the BCS and the IET.



**Citation on deposit:** Miao, X., Duan, H., Bai, Y., Shah, T., Song, J., Long, Y., Ranjan, R., & Shao, L. (online). Laser: Efficient Language-Guided Segmentation in Neural Radiance Fields. IEEE Transactions on Pattern Analysis and Machine Intelligence,

<https://doi.org/10.1109/TPAMI.2025.3535916>

**For final citation and metadata, visit Durham Research Online URL:**

<https://durham-repository.worktribe.com/output/3681375>

**Copyright statement:** This accepted manuscript is licensed under the Creative Commons Attribution 4.0 licence.

<https://creativecommons.org/licenses/by/4.0/>